A systematic optimization approach for a class of statistical inference problems utilizing data augmentation

Juan C. Agüero[†,‡], Rodrigo Carvajal[†] and Dimitrios Katselis.[*]

[†]Electronics Engineering Department, Universidad Técnica Federico Santa María, Chile.
✉ {juan.aguero,rodrigo.carvajalg}@usm.cl
[‡]School of Electrical Engineering and Computer Science, The University of Newcastle, Australia.
✉ juan.aguero@newcastle.edu.au
[*]Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, IL, USA. ✉ katselis@illinois.edu

arXiv:1601.05074v3 [math.OC] 15 Oct 2016

Abstract

We present an algorithm for a class of statistical inference problems. The main idea is to reformulate the inference problem as an optimization procedure, based on the generation of surrogate (auxiliary) functions. This approach is motivated by the MM algorithm, combined with the systematic and iterative structure of the Expectation-Maximization algorithm. The resulting algorithm can deal with hidden variables in Maximum Likelihood and Maximum a Posteriori estimation problems, Instrumental Variables, Regularized Optimization and Constrained Optimization problems.

The advantage of the proposed algorithm is to provide a systematic procedure to build surrogate functions for certain kind of systems typically arising in communication and quantization applications, where hidden variables are usually involved. Numerical examples show the benefits of the proposed approach.

1 Introduction

Problems in statistics and system identification often involve variables for which measurements are not available. Amongst others, real-life examples can be found in communication systems [5, 14] and systems with quantized data [13, 42]. For Maximum Likelihood (ML) estimation, the *likelihood function* is in general difficult to optimize by using closed-form expressions, and numerical approximations are usually cumbersome. These difficulties are traditionally avoided by the utilization of the Expectation-Maximization (EM) algorithm [7], where a surrogate (auxiliary) function is optimized instead of the main objective function. This surrogate function includes the complete data, i.e. the measurements and the random variables for which there are no measurements. The incorporation of such *hidden data* or *latent variables* is usually termed as *data augmentation*, where the main goal is to obtain, in general, simple and fast algorithms [41].

On the other hand, the MM[1] algorithm [22] is generally employed for solving more general optimization problems (not only for Maximum Likelihood and Maximum a Posteriori problems), particularly when there is no closed-form expressions for the solution of the optimization problem or the objective cost function is not convex. Real-life applications where the MM algorithm has been utilized include communication systems problems [32] and image processing [9]. For constrained optimization problems, an elegant solution is presented by Marks and Wright [29], where the constraints are incorporated via the formulation of surrogate functions. Surprisingly, Marks' approach has not received the same attention from the scientific community when it comes to compare it with the EM and the MM algorithms. In fact, these three approaches are contemporary, but the EM algorithm

[1]MM stands for Maximization-Minorization or Minimization-Majorization, depending on the optimization problem that needs to be solved.

has attracted most of the attention (out of the three methods), and it has been used for linear and nonlinear statistical inference in biology and engineering, see e.g. [1, 4, 16, 21, 35, 44], amongst others. On the other hand, as shown in [22], the MM algorithm has obtained much less attention, while Marks' approach is known only to a limited audience in the the Communication Systems community. These three approaches have important similarities: i) a surrogate function is defined and optimized in place of the original optimization problem, and ii) the solution is obtained iteratively. In general, these algorithms are "principles and recipes" [31] or a "philosophy" [22] for constructing solutions to a broad variety of optimization problems.

In this paper we adopt the ideas behind [7, 22, 29] to develop an algorithm for a special class of functions. Our approach generalizes the ones of [7, 22, 29] by expressing the cost function in terms of a kernel. In particular cases, the kernel corresponds to a variance-mean Gaussian mixture (VMGM), see e.g. [33]. VMGMs, also referred to as normal variance-mean mixtures [3] and normal scale mixtures [43], have been considered in the literature for formulating EM-based approaches to solve ML [2] and Maximum a Posteriori (MAP) problems [6, 12, 33]. Our approach is applicable to a wide class of functions, which allows for defining the likelihood function, the prior density function, and constraints as kernels, extending also the work in [29]. Thus, our work encompasses the following contributions: i) a systematic approach to constructing surrogate functions for a class of cost functions and constraints, ii) a class of kernels where the unknown quantities of the algorithm can be easily computed, and iii) a generalization of [7, 22, 29] by including the cost function and the constraints in one general expression.

Our algorithm is tested by an example, where we considered a system modelled by a linear regression. In this example the regression matrix is known and the parameter vector is unknown. The problem is solved considering a constrained optimization problem, where the (inequality) constraint is given by a multivariate ℓ_q-norm of the parameter vector. In this example, we show that the solution given by our proposal can iteratively converge to the global optimum.

2 Constrained problems in Statistical Inference

System Identification techniques encompass a variety of methodologies that can be used in order to obtain a model of a system from data. Classical methods, such as *Least Squares*, *Maximum Likelihood*, *Maximum a Posteriori* [15], *Prediction Error Method*, *Instrumental Variables* [36], and *Stochastic Embedding* [27] have been considered in the literature for such task. However, the increasing complexity of modern systems models have motivated researchers to revisit and reconsider those techniques for some problems. This has resulted in the incorporation of constraints and penalties into the identification problem, yielding an often more complicated optimization problem. For instance, it has been shown that the incorporation of linear equality constraints provide an improvement in the accuracy of the parameter estimates, see e.g [28]. On the other hand, the incorporation of regularization terms (or penalties) also improves the accuracy of the estimates, reducing the effect of noise and eliminating spurious local minima [25]. Regularization can be mainly incorporated in two ways: by adding regularizing constraints (a penalty function) or by including a probability density function (pdf) as a prior distribution for the parameters, see e.g. [12]. Another way to improve the estimation is by incorporating inequality constraints, where certain functions of the parameters may be required, for physical reasons amongst others, to be between certain bounds [20]. From this point of view, it is possible to consider the classical identification methods with constraints or penalties, as in [20, 23, 25, 27].

Perhaps one of the most utilized approaches for penalized estimation (with complicated non-linear expressions) is the MM algorithm – for details on the MM algorithm see Appendix A.2. This technique allows for the utilization of a surrogate function that is simple to handle, in terms of derivatives and optimization techniques, and that is, in turn, iteratively solved. However, its inequality constraint counterpart, here referred to as Marks' approach [29], is not as well known as the MM algorithm. Moreover, there is no straightforward manner to obtain such surrogate function. In this paper we focus on a systematic way to obtain the corresponding surrogate function using Marks' approach for a class of constraints.

3 Marks' approach for constrained optimization

The approach in [29] deals with inequality constraints by using a similar approach to the EM and MM algorithms. The basic idea is, again, to generate a surrogate function that allows for an iterative procedure whose optimum value is the optimum value of the original optimization problem.

Let us consider the following constrained optimization problem:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) \\ \text{s. t.} \quad & g(\boldsymbol{\theta}) \leq 0, \end{aligned} \tag{1}$$

where $f(\boldsymbol{\theta})$ is the objective function and $g(\boldsymbol{\theta})$ encodes the constraint of the optimization problem. In particular, let us focus on the case where $g(\boldsymbol{\theta})$ is not a convex function. This implies that the optimization problem cannot be directly solved. This difficulty can be overcome by utilizing a surrogate function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ at a given estimate $\hat{\boldsymbol{\theta}}^{(i)}$, such that

$$g(\boldsymbol{\theta}) \leq \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) \tag{2}$$

$$g(\hat{\boldsymbol{\theta}}^{(i)}) = \mathcal{Q}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) \tag{3}$$

$$\frac{d}{d\boldsymbol{\theta}} g(\boldsymbol{\theta})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = \frac{d}{d\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} \tag{4}$$

Provided the above properties are satisfied, then the following approximation of (1):

$$\begin{aligned} \boldsymbol{\theta}^{(i+1)} &= \arg\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) \\ \text{s. t.} \quad & \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) \leq 0, \end{aligned} \tag{5}$$

iteratively converges to the solution of the optimization problem (1). As shown in [29], the optimization problem in (5) is equivalent to the original problem (1), since the solution of (5) converges to a point that satisfies the Karush-Kuhn-Tucker conditions of the original optimization problem.

Remark 3.1 *Mark's approach can be considered as a generalization of the MM algorithm, since the latter can be derived (for a broad class of problems) from the former. Let us consider the following problem:*

$$\boldsymbol{\theta}^* = arg\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}). \tag{6}$$

Using the epigraph representation of (6) [19]*, we obtain the equivalent problem*

$$\begin{aligned} \boldsymbol{\theta}^* &= arg\min_{\boldsymbol{\theta}} t \\ s.\ t. \quad & f(\boldsymbol{\theta}) \leq t, \end{aligned} \tag{7}$$

Using Mark's approach (5)*, we can iteratively find a local optimum of* (6) *via*

$$\begin{aligned} \boldsymbol{\theta}^{(i+1)} &= arg\min_{\boldsymbol{\theta}} t \\ s.\ t. \quad & \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) \leq t, \end{aligned} \tag{8}$$

where $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ *in* (8) *is a surrogate function for* $f(\boldsymbol{\theta})$ *in* (6)*. From the epigraph representation we then obtain*

$$\boldsymbol{\theta}^{(i+1)} = arg\min_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}), \tag{9}$$

which is the definition of the MM algorithm (see Appendix A.2) for more details. ▽

4 A systematic approach to construct surrogate functions for a class of optimization problems

Here, we consider a general optimization cost defined as:

$$\mathcal{V}(\boldsymbol{\theta}) = \int\limits_{\Omega(\boldsymbol{y})} K(\boldsymbol{z}, \boldsymbol{\theta}) d\mu(\boldsymbol{z}), \tag{10}$$

where $\boldsymbol{\theta}$ is a parameter vector, $\boldsymbol{y}$ is a given data (i.e. measurements), $\boldsymbol{z}$ is the *complete data* (comprised of the *observed data* $\boldsymbol{y}$ and the *hidden variables* (unobserved data), $\Omega(\boldsymbol{y})$ is a mapping from the sample space of $\boldsymbol{z}$ to the sample space of $\boldsymbol{y}$, $K(\cdot,\cdot)$ is a (positive) kernel function, and $\mu(\cdot)$ is a measure, see e.g [8]. The definition of (10) is based on the definition of the auxiliary function $\mathcal{Q}$ in the EM algorithm [7, eq. (1.1)], where it is assumed throughout the paper that there is a mapping that relates the *not observed data* to the *observed data*, and that the *complete data* lies in $\Omega(\boldsymbol{y})$ [7]. Notice that in (10) the kernel function may not be a pdf. However, several functions can be expressed in terms of a pdf. The most common cases are Gaussian kernels (yielding VMGMs) [3] and Laplace kernels (yielding Laplace mixtures) [11].

Since we are considering the optimization of the function $\mathcal{V}(\boldsymbol{\theta})$, we can also consider the optimization of the function

$$\mathcal{J}(\boldsymbol{\theta}) = \log \mathcal{V}(\boldsymbol{\theta}). \tag{11}$$

From (11), we define the following auxiliary functions:

$$\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \int\limits_{\Omega(\boldsymbol{y})} \log[K(\boldsymbol{z}, \boldsymbol{\theta})] \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}), \tag{12}$$

$$\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \int\limits_{\Omega(\boldsymbol{y})} \log\left[\frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\mathcal{V}(\boldsymbol{\theta})}\right] \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}), \tag{13}$$

where $\hat{\boldsymbol{\theta}}^{(i)}$ is the ith estimate of $\boldsymbol{\theta}$. Thus, from (12) and (13) we obtain

$$\mathcal{J}(\boldsymbol{\theta}) = \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}). \tag{14}$$

As in the EM algorithm, for any $\boldsymbol{\theta}$, and using Jensen's inequality, it can be shown that the auxiliary function $\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ in (13) is a decreasing one.

Remark 4.1 *The kernel function $K(\boldsymbol{z}, \boldsymbol{\theta})$ satisfies the standing assumption $K(\boldsymbol{z}, \boldsymbol{\theta}) > 0$ since the proposed scheme is built, among others,on the logarithm of the kernel function $K(\boldsymbol{z}, \boldsymbol{\theta})$. The definition of the kernel function in* (10) *allows for kernels that are not probability density functions (pdf's). On the other hand, some kernels may correspond to a scaled version of a pdf. In that sense, for the cost function in* (10) *we can define a new kernel and a new measure as*

$$\bar{K}(\boldsymbol{z}, \boldsymbol{\theta}) = \frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\int K(\boldsymbol{z}, \boldsymbol{\theta}) d\boldsymbol{\theta}}, \; \bar{\mu}(\boldsymbol{z}) = \left(\int K(\boldsymbol{z}, \boldsymbol{\theta}) d\boldsymbol{\theta} \right) \mu(\boldsymbol{z}),$$

$$\Rightarrow \mathcal{V}(\boldsymbol{\theta}) = \int\limits_{\Omega(\boldsymbol{y})} \bar{K}(\boldsymbol{z}, \boldsymbol{\theta}) d\bar{\mu}(\boldsymbol{z}).$$

$\triangledown$

The idea behind using a surrogate function is to obtain a simpler algorithm for the optimization of the objective function when compared to the original optimization problem. This can be achieved iteratively if the Fisher Identity for the surrogate function and the objective function is satisfied. That is,

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = \frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}}. \tag{15}$$

Lemma 4.1 *For the class of objective functions in* (10)*, the surrogate function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ in* (12) *satisfies the Fisher identity defined in* (15)*.*

Table 1: Proposed algorithm.

Step 1:	Find a kernel that satisfies (10).
Step 2:	Obtain an initial guess $\hat{\boldsymbol{\theta}}^{(0)}$.
Step 3:	$i = 1$.
Step 4:	Compute $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ as in (12).
Step 5:	Compute $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$.
Step 6:	Incorporate $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ in the optimization problem and solve.
Step 7:	$i = i + 1$ and back to Step 4 until convergence.

Proof: From (14) we have:

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = \frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} - \frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}}.$$

Next, let us consider the gradient of the auxiliary function $\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$:

$$\begin{aligned}\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})\Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} &= \int_{\Omega(\boldsymbol{y})} \left[\frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\mathcal{V}(\boldsymbol{\theta})}\right]^{-1}_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} \frac{\partial}{\partial \boldsymbol{\theta}} \left[\frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\mathcal{V}(\boldsymbol{\theta}}\right]_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}) \\ &= \int_{\Omega(\boldsymbol{y})} \frac{\partial}{\partial \boldsymbol{\theta}} \left[\frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\mathcal{V}(\boldsymbol{\theta})}\right]_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} d\mu(\boldsymbol{z}) = \frac{\partial}{\partial \boldsymbol{\theta}} \left[\int_{\Omega(\boldsymbol{y})} \frac{K(\boldsymbol{z}, \boldsymbol{\theta})}{\mathcal{V}(\boldsymbol{\theta})} d\mu(\boldsymbol{z})\right]_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = 0.\end{aligned}$$

Hence, (15) holds. ■

Remark 4.2 *Note that the Fisher identity in Lemma 4.1 is well known in the EM-framework. However, we have specialized this result for the problem in this paper (i.e. when $K(\boldsymbol{z}, \boldsymbol{\theta})$ is not necessarily a probability density function)* ▽

Lemma 4.2 *The surrogate function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ in* (12) *can be utilized to obtain an adequate surrogate function that satisfies the properties in Mark's approach in* (2)-(4).

Proof: Notice that in Mark's approach the optimization problem corresponds to the minimization of the objective function. Hence, to maximize, we have $\mathcal{J}(\boldsymbol{\theta}) = -\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) + \mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$. From (14) we can construct the surrogate functions $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ and $\tilde{\mathcal{H}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ since

$$\mathcal{J}(\boldsymbol{\theta}) = -\underbrace{(\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \mathcal{Q}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) + \mathcal{J}(\hat{\boldsymbol{\theta}}^{(i)}))}_{\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})} + \underbrace{(\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \mathcal{Q}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) + \mathcal{J}(\hat{\boldsymbol{\theta}}^{(i)}))}_{\tilde{\mathcal{H}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})}. \tag{16}$$

The function $\tilde{\mathcal{H}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ satisfies $\tilde{\mathcal{H}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \tilde{\mathcal{H}}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) \geq 0$, which implies that $\mathcal{J}(\boldsymbol{\theta}) \leq \tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$, satisfying (2). From $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \mathcal{Q}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) + \mathcal{J}(\hat{\boldsymbol{\theta}}^{(i)}$ we obtain $\tilde{\mathcal{Q}}(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) = \mathcal{J}(\hat{\boldsymbol{\theta}}^{(i)})$, satisfying (3). Finally, since $\tilde{\mathcal{H}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ satisfies (15), $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ satisfies (4). ■

Remark 4.3 *Since $\frac{d}{d\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \frac{d}{d\boldsymbol{\theta}} \tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$, it is simpler to consider the function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ instead of $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ in penalized/regularized and MAP estimation problems, as shown in* [12] *and* [6]. ▽

In Table 1 we summarize our proposed algorithm.

5 A quadratic surrogate function for a class of kernels

In this paper we focus on a special class of the kernel functions $K(\boldsymbol{z}, \boldsymbol{\theta})$. For this particular class, the following is satisfied:

$$\frac{\partial}{\partial \boldsymbol{\theta}} \log[K(\boldsymbol{z}, \boldsymbol{\theta})] = \mathbf{A}(\boldsymbol{z})\boldsymbol{\theta} + \mathbf{b}, \tag{17}$$

where $\mathbf{A}(\boldsymbol{z})$ is a matrix and $\mathbf{b}$ is a vector, both of adequate dimensions. Then, we have that

$$\begin{aligned}
\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) &= \int_{\Omega(\boldsymbol{y})} [\mathbf{A}(\boldsymbol{z})\boldsymbol{\theta} + \mathbf{b}] \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}) \\
&= \left[\int_{\Omega(\boldsymbol{y})} \mathbf{A}(\boldsymbol{z}) \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}) \right] \boldsymbol{\theta} + \mathbf{b} \int_{\Omega(\boldsymbol{y})} \frac{K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)})}{\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})} d\mu(\boldsymbol{z}) \\
&= \mathbf{R}\boldsymbol{\theta} + \mathbf{b}.
\end{aligned} \tag{18}$$

Remark 5.1 *Notice that in the previous expression is linear with respect to $\boldsymbol{\theta}$. Hence, the surrogate function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ is quadratic with respect to the parameter vector $\boldsymbol{\theta}$.* ▽

From the Fisher Identity in (15) we have that

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = \mathbf{R}\hat{\boldsymbol{\theta}}^{(i)} + \mathbf{b}, \tag{19}$$

from which we can solve for $\mathbf{R}$ in some cases[2]. In particular, if $\mathbf{A}(\boldsymbol{z})$ is a diagonal matrix, then $\mathbf{R}$ is also a diagonal matrix defined by $\mathbf{R} = \text{diag}[r_1, r_2, ...]$. Thus, we have

$$\frac{\partial}{\partial \theta_i} \mathcal{J}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} = r_i \hat{\theta}_i^{(i)} + b_i \Rightarrow r_i = \frac{\frac{\partial}{\partial \theta_i} \mathcal{J}(\boldsymbol{\theta}) \Big|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}^{(i)}} - b_i}{\hat{\theta}_i^{(i)}},$$

where θ_i is the ith component of the parameter vector $\boldsymbol{\theta}$, $\hat{\boldsymbol{\theta}}_i^{(i)}$ is the ith component of the estimate $\hat{\boldsymbol{\theta}}^{(i)}$, r_i is the ith element of the diagonal of $\mathbf{R}$, and b_i is the ith element of the vector $\mathbf{b}$. Hence, when optimizing the auxiliary function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ we obtain

$$\frac{\partial}{\partial \boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \begin{bmatrix} r_1 & & \\ & r_2 & \\ & & \ddots \end{bmatrix} \boldsymbol{\theta} + \mathbf{b} = 0 \Rightarrow \hat{\theta}_i^{(i+1)} = -\frac{b_i}{r_i}.$$

Equivalently,

$$\hat{\boldsymbol{\theta}}^{(i+1)} = \mathbf{R}^{-1}\mathbf{b}. \tag{20}$$

This implies that in our approach, it is not necessary to obtain the auxiliary function $\mathcal{Q}$ and optimize it. Instead, by computing $\mathbf{R}$ and $\mathbf{b}$, the new estimate can be obtained.

For the class of kernels here described, the proposed method for constructing surrogate functions can also be understood as part of sequential quadratic programming (SQP) methods [10, ch. 12.4]. Indeed, the general case of equality and inequality-constrained minimization problems is defined as [24, ch. 4]:

$$\begin{aligned}
\boldsymbol{\theta}^* &= \arg\min_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) \\
&\text{s. t.} \quad h(\boldsymbol{\theta}) = 0\,,\; g(\boldsymbol{\theta}) \leq 0,
\end{aligned}$$

which is solved by iteratively defining quadratic functions that approximate the objective function and the inequality constraint around a current iterate $\hat{\boldsymbol{\theta}}^{(i)}$. In the same way, our proposal generates an algorithm with quadratic surrogate functions.

[2]The matrix $\mathbf{R}$ can also be computed using Monte Carlo algorithms.

Remark 5.2 *In the proposed methodology, the surrogate function that can be optimized is given by*

$$\bar{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \int_{\Omega(\boldsymbol{y})} \log[K(\boldsymbol{z}, \boldsymbol{\theta})] K(\boldsymbol{z}, \hat{\boldsymbol{\theta}}^{(i)}) d\mu(\boldsymbol{z}), \tag{21}$$

since $\mathcal{V}(\hat{\boldsymbol{\theta}}^{(i)})$ *in* (12) *does not depend on the parameter* $\boldsymbol{\theta}$*. Thus, the proposed method corresponds to a variation of the EM algorithm that is not limited to probability density functions (e.g. the likelihood function) for solving ML and MAP estimation problems. Instead, our version considers general measures* ($\mu(\boldsymbol{z})$)*, where the mapping over the measurement data* $\Omega(\boldsymbol{y})$ *is a given set.* ▽

6 Numerical Examples

In this section, we illustrate our proposed algorithm with the following inequality constrained ML estimation problem:

$$\begin{aligned} \boldsymbol{\theta}^* &= \arg\min_{\boldsymbol{\theta}} \|\boldsymbol{y} - \boldsymbol{A}\boldsymbol{\theta}\|_2^2 \\ \text{s. t. } f(\boldsymbol{\theta}) &= \sum_{m=1}^{M} \left\| \frac{\sqrt{\beta_m}}{\tau} \boldsymbol{\theta}_m \right\|_2^q \leq \gamma, \end{aligned} \tag{22}$$

where $\boldsymbol{\theta} = [\boldsymbol{\theta}_1^T, ..., \boldsymbol{\theta}_M^T]^T$ is an unknown vector parameter, $\boldsymbol{A}$ is the regressor matrix, and $\boldsymbol{y}$ corresponds to the measurements, with $q = 0.4$, $\tau = 0.2$, and $\gamma = 7$. Notice that the constraint function belongs to a family of sparsity constraints found in sparse system identification and optimization – see e.g [23] and the references therein. The function of the unknown parameter in the inequality constraint can be understood as its "group ℓ_q-norm", with M groups and where $\boldsymbol{\theta}_m$ is the mth group of length ${\beta_m}$[3]. Hence, the inequality constraint can be expressed as a Multivariate Power Exponential (MPE) distribution [37] of $\boldsymbol{\theta}$. On the other hand, the MPE distributions can be expressed as VMGMs [26], which allows for utilizing the ideas presented in Section 4, where the corresponding kernel function is a multivariate Gaussian distribution.The elements of the matrix $\boldsymbol{A}$ were generated using a Normal distribution ($\mathcal{N}(0, 1)$) and the measurement data was generated from

$$\boldsymbol{y} = \boldsymbol{A}\boldsymbol{\theta} + \boldsymbol{n},$$

where $\boldsymbol{n} \sim \mathcal{N}(\boldsymbol{0}, 0.01\mathbf{I})$ is the additive measurement noise. In this example, we consider that the number of measurements is 256 and that $\boldsymbol{\theta} = [0.54, 1.83, -2.26, 0, 0,\ 0, 0, 0, 0, 0.86, 0.32, -1.31]^T$, with $\beta_m = 3\ \forall m$, and thus $M = 4$. We have also considered 150 Monte Carlo simulations.

For the attainment of the surrogate function we consider the following:

1. When expressed as a MPE distribution, the group ℓ_q-norm is given by [37],

$$p(\boldsymbol{\theta}_m) = \kappa_m(q, \tau) e^{-\left\| \frac{\sqrt{\beta_m}}{\tau} \boldsymbol{\theta}_m \right\|_2^q},$$

   where $\kappa_m(q, \tau)$ is a constant that depends on τ and q, and

$$p(\boldsymbol{\theta}) = \prod_{m=1}^{M} p(\boldsymbol{\theta}_m).$$

2. When expressed as a VMGM, the group ℓ_q-norm is given by

$$p(\boldsymbol{\theta}_m) = \int \mathcal{N}_{\boldsymbol{\theta}}(0, \boldsymbol{\Sigma}(\lambda_m)) p(\lambda_m) d\lambda_m.$$

[3]Notice that when $m = 1$ we obtain the standard ℓ_q-norm, and when $q = 1$ we obtain the ℓ_1-norm utilized in the Lasso.

3. With the previous expressions, and using the results in Section 5, we notice that

$$\frac{\partial \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})}{\boldsymbol{\theta}_m} = -k_m^{(i)} \boldsymbol{\theta}_m,$$

where $\hat{\boldsymbol{\theta}}_m^{(i)}$ is the estimate of the mth group at the ith iteration and $k_m^{(i)}$ is a constant that depends on $\hat{\boldsymbol{\theta}}_m^{(i)}$.

4. Then we obtain $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \frac{k_m^{(i)}}{2} \left(\boldsymbol{\theta}_m^T \boldsymbol{\theta}_m \right) + c$, where c is a constant.

Then, the inequality constraint in (22) can be replaced with the following constraint (that includes the surrogate function $\tilde{\mathcal{Q}}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ as shown in (16))

$$\sum_{m=1}^{M} \frac{k_m^{(i)}}{2} \left(\boldsymbol{\theta}_m^T \boldsymbol{\theta}_m - (\hat{\boldsymbol{\theta}}_m^{(i)})^T (\hat{\boldsymbol{\theta}}_m^{(i)}) \right) + \left\| \frac{\sqrt{\beta_m}}{\tau} \hat{\boldsymbol{\theta}}_m^{(i)} \right\|_2^q \leq \gamma, \tag{23}$$

The results of the optimization problem for one realization is shown in Fig. 1, where the optimization problem with the surrogate function was solved using the optimization software CVX [18, 17] and the original non-convex problem was solved using the optimization software BARON [39, 34], which allowed us to obtain the global optimum. From Fig. 1 it can be seen that our approach finds the global minimum with a few iterations only. On the other hand, we have also considered the mean squared error (MSE) of the estimates and compared them with the unconstrained problem (least squares). The MSE obtained with our approach is 3.6×10^{-2}, whilst the least squares estimate yielded an MSE equal to 4.34×10^{-2}. Clearly, the incorporation of the constraint aided the estimation.

7 Conclusions

In this paper we have presented a systematic approach for constructing surrogate functions in a wide range of optimization problems, generalizing the popular EM and MM algorithms. Our approach is based on the utilization of data augmentation and kernel functions, yielding very simple optimization algorithms when the kernel can be expressed as VMGM.

A Rudiments of the proposed approach

A.1 The EM algorithm

Let us consider an estimation problem and its corresponding log-likelihood function defined as $\ell(\boldsymbol{\theta}) = \log p(\boldsymbol{y}|\boldsymbol{\theta})$, where $p(\boldsymbol{y}|\boldsymbol{\theta})$ is the likelihood function, $\boldsymbol{\theta} \in \mathbb{R}^p$, and $\boldsymbol{y} \in \mathbb{R}^N$. Denoting the *complete data* by $\boldsymbol{z} \in \Omega(\boldsymbol{y})$, and using Bayes' theorem, we can obtain:

$$\ell(\boldsymbol{\theta}) = \log p(\boldsymbol{y}|\boldsymbol{\theta}) = \log p(\boldsymbol{z}|\boldsymbol{\theta}) - \log p(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta}). \tag{24}$$

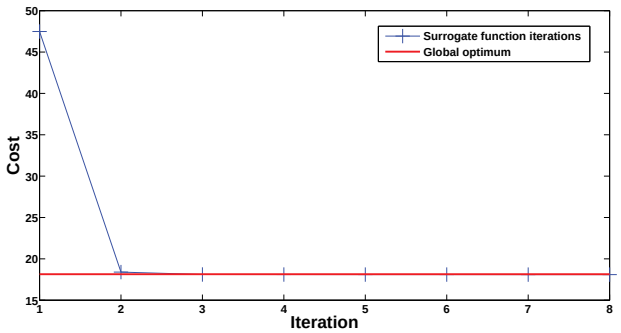

Figure 1: Convergence of the proposed approach to the global optimum

Let us assume that at the ith iteration we have the estimate $\hat{\boldsymbol{\theta}}^{(i)}$. By integrating at both sides of (24) with respect to $p(\boldsymbol{z}|\boldsymbol{y}, \hat{\boldsymbol{\theta}}^{(i)})$ we obtain $\ell(\boldsymbol{\theta}) = \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) - \mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$, where

$$\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \int_{\Omega(\boldsymbol{y})} \log p(\boldsymbol{z}|\boldsymbol{\theta}) p(\boldsymbol{z}|\boldsymbol{y}, \hat{\boldsymbol{\theta}}^{(i)}) d\boldsymbol{z}, \tag{25}$$

$$\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = \int_{\Omega(\boldsymbol{y})} \log p(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta}) p(\boldsymbol{z}|\boldsymbol{y}, \hat{\boldsymbol{\theta}}^{(i)}) d\boldsymbol{z}. \tag{26}$$

Using Jensen's inequality [8, pp. 24–25], it is possible to show that for any value of $\boldsymbol{\theta}$, the function $\mathcal{H}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ is decreasing. Hence, the optimization is only carried out on the auxiliary function $\mathcal{Q}$ because, by maximizing $\mathcal{Q}$, the new parameter $\hat{\boldsymbol{\theta}}^{(i+1)}$ is such that the likelihood function increases (see e.g. [7, 30]).

In general, the EM method can be summarised as follows:

E-step: Compute the expected value of the joint *likelihood function* for the *complete data* (measurements and hidden variables) based on a given parameter estimate, $\hat{\boldsymbol{\theta}}^{(i)}$. Thus, we have (see e.g. [7]):

$$\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) = E[\log p(\boldsymbol{z}|\boldsymbol{\theta})|\boldsymbol{y}, \hat{\boldsymbol{\theta}}^{(i)}], \tag{27}$$

M-step: Maximize the function $\mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)})$ (27), with respect to $\boldsymbol{\theta}$:

$$\hat{\boldsymbol{\theta}}^{(i+1)} = \arg\max_{\boldsymbol{\theta}} \mathcal{Q}(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}). \tag{28}$$

This succession of estimates converges to a stationary point of the *log-likelihood* function [40].

A.2 The MM algorithm

The idea behind the MM algorithm [22] is to construct a surrogate function $g(\boldsymbol{\theta}, \boldsymbol{z})$, that majorizes (for minimization problems) or minorizes (for maximization problems) a given cost functions $f(\boldsymbol{\theta})$ [22] such that,

$$\begin{aligned} &f(\boldsymbol{\theta}) \leq g(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) && \text{for minimization problems, or} \\ &f(\boldsymbol{\theta}) \geq g(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}) && \text{for maximization problems, and} \\ &f(\boldsymbol{\theta}) = g(\boldsymbol{\theta}, \boldsymbol{\theta}), \end{aligned}$$

where $\hat{\boldsymbol{\theta}}^{(i)}$ is an estimate of $\boldsymbol{\theta}$. Then, the surrogate function is iteratively optimized until convergence. Hence, for maximizing $f(\boldsymbol{\theta})$ we have [38]

$$\hat{\boldsymbol{\theta}}^{(i+1)} = \arg\max_{\boldsymbol{\theta}} g(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}}^{(i)}). \tag{29}$$

For the construction of the surrogate function, popular techniques include the first order Taylor approximation, the quadratic upper bound principle and Jensen's inequality for convex functions, see, e.g., [38].

Remark A.1 *The iterative strategy utilized in the MM algorithm converges to a local optimum since*

$$f(\hat{\boldsymbol{\theta}}^{(i+1)}) \geq g(\hat{\boldsymbol{\theta}}^{(i+1)}, \hat{\boldsymbol{\theta}}^{(i)}) \geq g(\hat{\boldsymbol{\theta}}^{(i)}, \hat{\boldsymbol{\theta}}^{(i)}) = f(\hat{\boldsymbol{\theta}}^{(i)}). \qquad \triangledown$$

Acknowledgements

The work of J. C. Agüero is supported by the National Fund for Scientific and Technological Development (FONDECYT) - Chile through through grant no. 1150954. The work of R. Carvajal is supported by FONDECYT - Chile through its Postdoctoral Fellowship Program 2014 – grant No. 3140054. This work was also partially supported by the Advanced Center for Electrical and Electronic Engineering (AC3E, Proyecto Basal FB0008), Chile.

References

[1] J. C. Agüero, W. Tang, J. I. Yuz, R. Delgado, and G. C. Goodwin. Dual time-frequency domain system identification. *Automatica*, 48(12):3031–3041, 2012.

[2] N. Balakrishnam, V. Leiva, A. Sanhueza, and F. Vilca. Estimation in the Birnbaum-Saunders distribution based on scale-mixture of normals and the EM algorithm. *SORT*, 33(2):171–192, 2009.

[3] O. Barndorff-Nielsen, J. Kent, and M. Sorensen. Normal variance-mean mixtures and z distributions. *Int. Stat. Review*, 50(2):145–159, 1982.

[4] M. J. Beal and Z. Ghahramani. The variational bayesian EM algorithm for incomplete data: with application to scoring graphical model structures. In *Proc. of the 7th Valencia International Meeting)*, Valencia, Spain, 2003.

[5] R. Carvajal, J. C. Agüero, B. I. Godoy, and G. C. Goodwin. EM-based Maximum-Likelihood channel estimation in multicarrier systems with phase distortion. *IEEE Trans. Vehicular Technol.*, 62(1):152–160, 2013.

[6] R. Carvajal, J. C. Agüero, B. I. Godoy, and D. Katselis. A MAP approach for ℓ_q-norm regularized sparse parameter estimation using the EM algorithm. In *Proc. of the 25th IEEE Int. Workshop on Mach. Learning for Signal Process (MLSP 2015)*, Boston, USA, 2015.

[7] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from imcomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B*, 39(1):1–38, 1977.

[8] R. Durret. *Probability: Theory and examples.* Cambridge University Press, 4th edition, 2010.

[9] M. A. T. Figueiredo, J. M. Bioucas-Dias, and R. D. Nowak. Majorization–Minimization algorithms for wavelet-based image restoration. *IEEE Trans. Image Process.*, 16(12):2980–2991, 2007.

[10] R. Fletcher. *Practical methods of optimization.* John Wiley & Sons, Chichester, GB, 2nd edition, 1987.

[11] P Garrigues and B. Olshausen. Group sparse coding with a laplacian scale mixture prior. *Advances in Neural Information Processing Systems*, 23:676–684, 2010.

[12] B. I. Godoy, J. C. Agüero, R. Carvajal, G. C. Goodwin, and J. I. Yuz. Identification of sparse FIR systems using a general quantisation scheme. *Int. J. Control*, 87(4):874–886, 2014.

[13] B. I. Godoy, G. C. Goodwin, J. C. Agüero, D. Marelli, and T. Wigren. On identification of FIR systems having quantized output data. *Automatica*, 47(9):1905–1915, 2011.

[14] A. Goldsmith. *Wireless Communications.* Cambridge University Press, New York, NY, USA, 2005.

[15] G. C. Goodwin and R. L Payne. *Dynamic system identification : experiment design and data analysis.* Academic Press New York, 1977.

[16] R. B. Gopaluni. A particle filter approach to identification of nonlinear processes under missing observations. *Can J. Chem. Eng.*, 86(6):1081–1092, 2008.

[17] M. Grant and S. Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008. `http://stanford.edu/~boyd/graph_dcp.html`.

[18] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014.

[19] M. C. Grant and S. P. Boyd. *Recent Advances in Learning and Control*, chapter Graph Implementations for Nonsmooth Convex Programs, pages 95–110. Springer, London, 2008.

[20] M. A. Hanson. Inequality constrained maximum likelihood estimation. *Ann Inst Stat Math*, 17(1):311–321, 1965.

[21] A. Hobolth and J. L. Jensen. Statistical inference in evolutionary models of DNA sequences via the EM algorithm. *Statistical Applications in Genetics and Molecular Biology*, 4(1), 2005. Article 18.

[22] D.R. Hunter and K. Lange. A tutorial on MM algorithms. *The American Statistician*, 58(1):30–37, 2004.

[23] M.M. Hyder and K. Mahata. A robust algorithm for joint-sparse recovery. *Signal Process. Letters, IEEE*, 16(12):1091–1094, Dec 2009.

[24] A. F. Izmailov and M. V. Solodov. *Newton-typo methods for optimization and variational problems*. Springer, Cham, Switzerland, 2014.

[25] K. Lange, editor. *Optimization*. Springer, New York, USA, 2nd edition, 2013.

[26] K. Lange and J. S. Sinsheimer. Normal/independent distributions and their applications in robust regression. *J. Comput. Graphical Stat.*, 2(2):175–198, 1993.

[27] L. Ljung, G. C. Goodwin, and J. C. Agüero. Stochastic embedding revisited: A modern interpretation. In *Proc. of the 53rd IEEE Conf. Decision and Control)*, Los Angeles, CA, USA, 2014.

[28] K. Mahata and T Söderström. Improved estimation performance using known linear constraints. *Automatica*, 40(8):1307 – 1318, 2004.

[29] B. R. Marks and G. P. Wright. A general inner approximation algorithm for nonconvex mathematical programs. *Operations Research*, 26(4):681–683, 1978.

[30] G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. Wiley, 1997.

[31] X. L. Meng. Thirty years of EM and much more. *Statistica Sinica*, 17(3):839–840, 2007.

[32] E. Pollakis, R. L. G. Cavalcante, and S. Stańczak. Base station selection for energy efficient network operation with the Majorization- Minimization algorithm. In *Proc. 13th IEEE Int. Workshop on Signal Process. Adv. Wireless Commun. (SPAWC 2012)*, Çeşme, Turkey, 2012.

[33] N.G. Polson and J.G. Scott. Data augmentation for non-gaussian regression models using variance-mean mixtures. *Biometrika*, 100:459–471, 2013.

[34] N. V. Sahinidis. *BARON 14.3.1: Global Optimization of Mixed-Integer Nonlinear Programs,* User's Manual, 2014.

[35] T. B. Schön, A. Wills, and B. Ninness. System identification of nonlinear state-space models. *Automatica*, 47(1):39–49, 2011.

[36] T. Söderström and P. Stoica, editors. *System Identification*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1988.

[37] N. Solaro. Random variate generation from multivariate exponential power distribution. *Statistica & Applicazioni*, 2(2):25–44, 2004.

[38] B. K. Sriperumbudur, D. A. Torres, and G. R. G. Lanckriet. A majorization-minimization approach to the sparse generalized eigenvalue problem. *Machine Learning*, 85(1–2):3–39, 2011.

[39] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103:225–249, 2005.

[40] F. Vaida. Parameter convergence for EM and MM algorithms. *Statistica Sinica*, 15(3):831–840, 2005.

[41] D. A. van Dyk and X.-L. Meng. The art of data augmentation. *Journal of Computational and Graphical Statistics*, 10(1):1–50, 2001.

[42] L. Wang, J. Zhang, and G. G. Yin. System identification using binary sensors. *IEEE Trans. Autom. Control*, 48(11):1892–1907, 2003.

[43] M. West. On scale mixtures of normal distributions. *Biometrika*, 74(3):646–648, 1987.

[44] S. Yang, J. K. Kim, and Z. Zhu. Parametric fractional imputation for mixed models with nonignorable missing data. *Statistics and Its Interface*, 6(3):339–347, 2013.